ARMY RESEARCH LABORATORY

ARL

U.S. Army Research Laboratory Image Enhancement Test Bed User's Manual

by Laurel C. Sadler

ARL-TR-6512   **July 2013**

NOTICES

Disclaimers

The findings in this report are not to be construed as an official Department of the Army position unless so designated by other authorized documents.

Citation of manufacturer's or trade names does not constitute an official endorsement or approval of the use thereof.

Destroy this report when it is no longer needed. Do not return it to the originator.

Army Research Laboratory

Adelphi, MD 20783-1197

ARL-TR-6512 **July 2013**

U.S. Army Research Laboratory Image Enhancement Test Bed User's Manual

Laurel C. Sadler
Computational and Information Sciences Directorate, ARL

REPORT DOCUMENTATION PAGE

1. REPORT DATE *(DD-MM-YYYY)*	2. REPORT TYPE	3. DATES COVERED (From - To)
July 2013	Final	October 2013 to April 2013

4. TITLE AND SUBTITLE	5a. CONTRACT NUMBER
U.S. Army Research Laboratory Image Enhancement Test Bed User's Manual	
	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)	5d. PROJECT NUMBER
Laurel C. Sadler	R.0006163.9
	5e. TASK NUMBER
	5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)	8. PERFORMING ORGANIZATION REPORT NUMBER
U.S. Army Research Laboratory ATTN: RDRL-CII-B 2800 Powder Mill Road Adelphi MD 20783-1197	ARL-TR-6512

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR/MONITOR'S ACRONYM(S)
	11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT

Approved for public release; distribution unlimited.

13. SUPPLEMENTARY NOTES
Email: laurel.c.sadler.civ@mail.mil

14. ABSTRACT

The U.S. Army Research Laboratory (ARL) Image Enhancement Test Bed was developed to evaluate various image processing research algorithms being developed at ARL. The Test Bed integrates these algorithms into a more real-world environment simulating Army-relevant scenarios. Some of the algorithms are made into client applications such as Super Resolution, Contrast, and Deblur, which can be accessed from an Image Processing Representational State Transfer (REST) Web service by a cell phone, laptop PC, Microsoft Surface Touch Table, as well as other devices. Other more central processing unit (CPU) intensive and time-constrained algorithms such as tracking algorithms run directly on the laptop PC or Touch Table Computer. This report briefly describes the image processing algorithms implemented in the ARL Image Enhancement Test Bed while citing the supporting research publications. The main purpose of this report is to provide a detailed description of the ARL Image Enhancement Test Bed Environment and graphical user interface (GUI), and serve as a user's manual.

15. SUBJECT TERMS

Image processing, super resolution, deblur, contrast, tracker

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE	UU	44	Laurel C. Sadler
Unclassified	Unclassified	Unclassified			19b. TELEPHONE NUMBER *(Include area code)* (301) 394-1221

Contents

List of Figures

1. Introduction

The U.S. Army Research Laboratory (ARL) Image Enhancement Test Bed was developed to evaluate various image processing research algorithms being developed at ARL. The Test Bed integrates these algorithms into a more real-world environment simulating Army-relevant scenarios. Some of the algorithms are made into client applications such as Super Resolution, Contrast, and Deblur, which can be accessed from an Image Processing Representational State Transfer (REST) Web service (*1*) by a cell phone, laptop PC, Microsoft Surface Touch Table, as well as other devices. Other more central processing unit (CPU) intensive and time-constrained algorithms such as tracking algorithms run directly on the laptop PC or Touch Table Computer. This report briefly describes the image processing algorithms implemented in the ARL Image Enhancement Test Bed while citing the supporting research publications. The main purpose of this report is to provide a detailed description of the ARL Image Enhancement Test Bed Environment and graphical user interface (GUI), and serve as a user's manual.

2. Capabilities and Specifications

The ARL Image Enhancement Test Bed GUI integrates the ARL image processing research algorithms into one control center. The GUI was designed to allow the operator the flexibility to run the image processing algorithms in a real-world environment, providing a streamlined, dynamic, reconfigurable workflow. The GUI allows for easy modification of the parameters for each of the algorithms; the ability to pan, tilt, and zoom the networked camera, providing the full motion video; as well as annotate and save the enhanced images for later analysis. The user interface and system control software were developed in .NET Framework 3.0 with Windows Presentation Foundation (WPF) using Microsoft Visual Studio (Version 10.0). The ARL Image Enhancement Test Bed application can be executed on both a laptop PC using the mouse and keyboard for input or the Microsoft Surface Touch Table. The ARL Image Enhancement Test Bed GUI is described in detail in section 3.

3. The ARL Image Enhancement Test Bed Graphical User Interface (GUI)

The main window for the ARL Image Enhancement GUI is shown in figure 1. From this window the operator can select from images stored on the Web service; apply three unique image enhancement algorithms to these images; annotate and save the outputs of the algorithms; view full motion video from the selected network camera; pan, tilt, and zoom the network camera;

apply the super resolution algorithm to the live video; and execute an image tracking algorithm on the full motion video. The operator can also view previously enhanced and annotated images that were saved to the Web service.
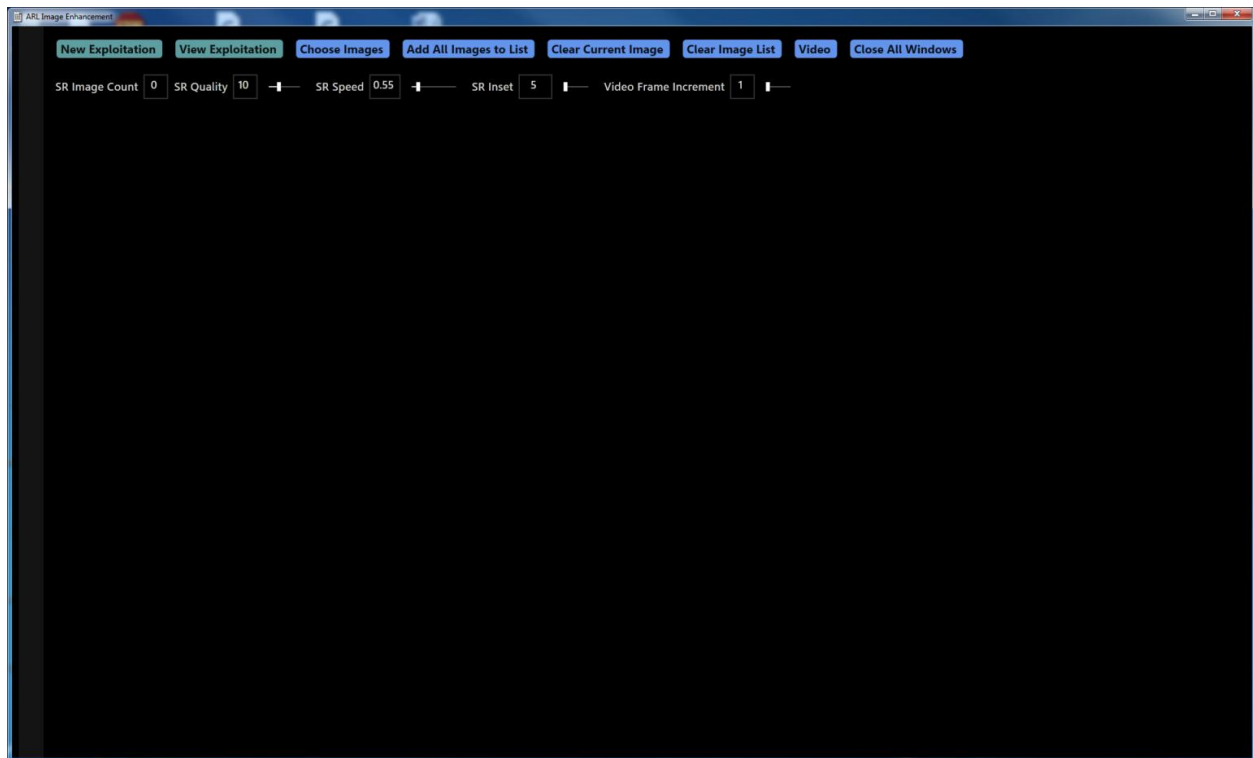
Figure 1. ARL Image Enhancement GUI initial screen.

When the blue Choose Images button at the top of the main screen is pressed, the Select a Collection of Images Pop Up window appears on the main screen, as shown in figure 2. From this menu, the operator can select from collections of images that have been stored on the Web service. The operator can click or touch the collection item, which then becomes highlighted.
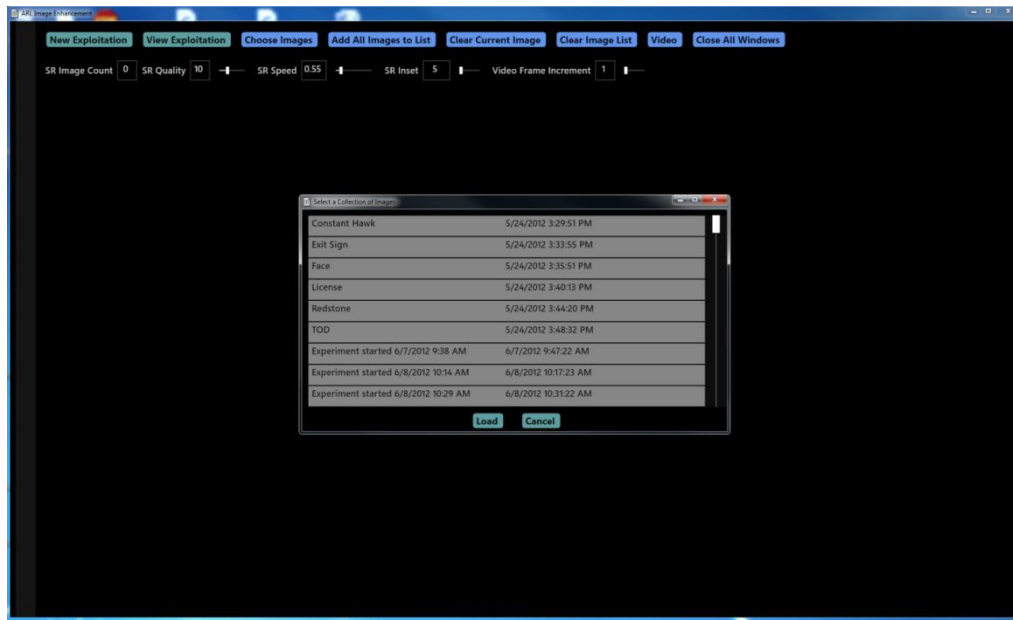
Figure 2. Select a Collection of Images from the REST Web service.

When the Load button is pressed, the images are retrieved from the Web service and displayed in a scrolling image list box on the left side of the GUI, as shown in figure 3. The operator has the option to scroll through the images selecting one image at a time by touching the image when using the Microsoft Surface or clicking the left mouse button on the image. These chosen images are displayed in the image stack box with the light steel blue background. The top image in the image stack box is the image currently displayed in the Image Enhancement window. The operator can also use the Add All Images to List button on the main screen to add all the images in the scrolling image list to the image list and image stack box. The Clear Image List button on the main screen clears all images from the image list and stack box. The Clear Current Image button removes the top most image in the stack box from the image list and the stack box. The images in the image list and stack box are used when super resolution is applied to an operator specified region of the image.
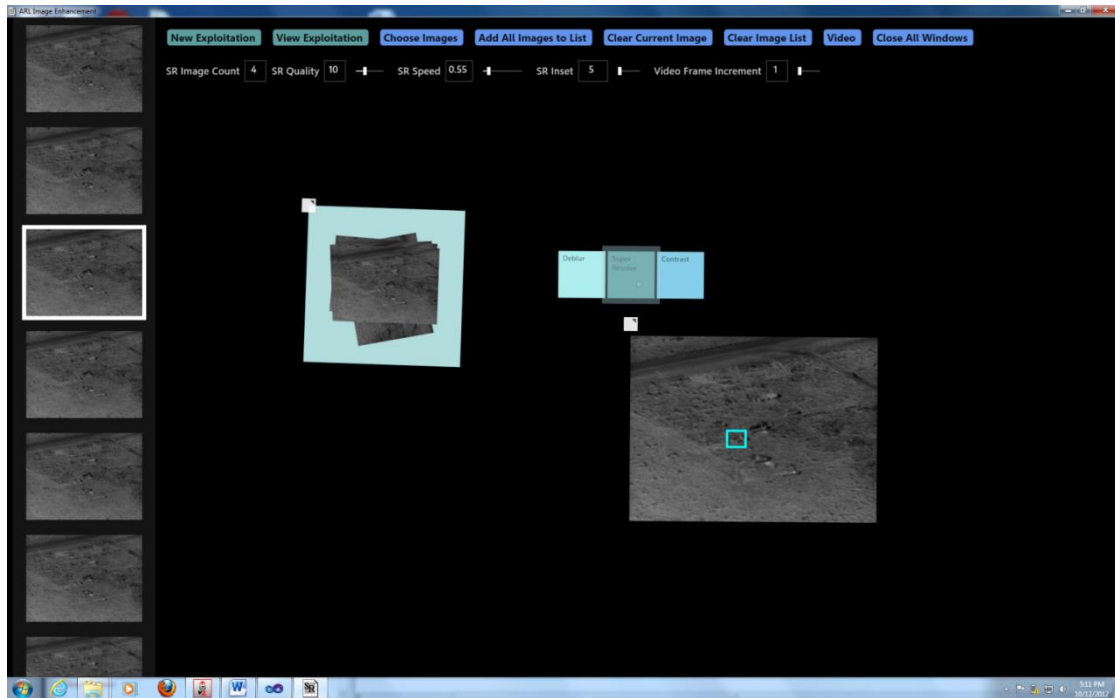
Figure 3. Collection of Images.

The SR Image Count displays the number of images currently in the stack box. Super resolution can only be applied to multiple images. If only one image is chosen when the super resolution client application is executed, the operator is notified with an error message, which is displayed at the bottom of the main window, as shown in figure 4. This message fades out in 30 s. If more than 20 images are selected when the super resolution client application is executed, the operator is notified with the error message, as shown in figure 5.
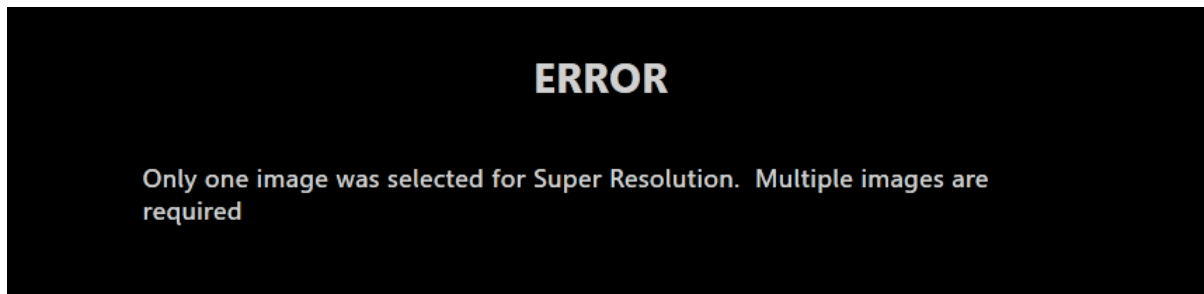
Figure 4. Error- Multiple images are required.
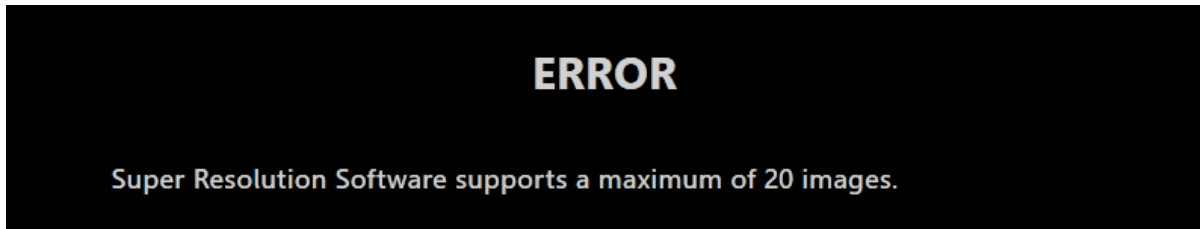
Figure 5. Error – Exceeded the maximum number of images allowed.

3.1   Super Resolution

The super resolution (SR) (*2, 3*) image reconstruction algorithm uses a sequence of low-resolution, under-sampled frames containing sub-pixel shift to generate a super-resolved image recovering high frequency information. Three parameters are used by the ARL SR algorithm: quality, speed, and inset. The main screen displays three sliders to allow the operator to adjust these parameters. SR quality is the up sampling factor for the reconstruction grid. It is an even integer and is typically set to 1.5 times the number of frames. The slider allows the operator to select even integers between 4 and 24. However, increasing quality lengthens runtime; thus, for quick demonstrations, a setting of 10 is appropriate. The SR speed controls the frequency bandwidth of the super-resolved image. It is a decimal between 0 and 1, and can typically be set to 0.7. Usually SR speed would be increased if the number of frames is small (<9 frames). The slider allows the user to set the SR speed from 0.5 to 1 in increments of 0.05. The last variable, inset, is an integer pixel inset variable that defines a sub-region centered within the user-defined region (designated by the blue rectangle) for frame registration purposes. A setting of 5 for inset will work in most scenarios. The slider allows the operator to set the SR inset to integers between 4 and 1000.

Before the SR client application can be executed the operator must select multiple images and designate the region of the image on which to apply the algorithm. The region is selected by moving and resizing the blue rectangle in the Image Enhancement window over the region to be super-resolved, as shown in figures 3 and 6. After the region is selected, the operator may move through the stack list of images by clicking on the top image of the stack box. This action also displays the top most image in the Image Enhancement window. Viewing each image in the Image Enhancement window allows the operator ensure that the selected region is correct for each image. The rectangle must be large enough to surround the entire region for enhancement on all of the images.
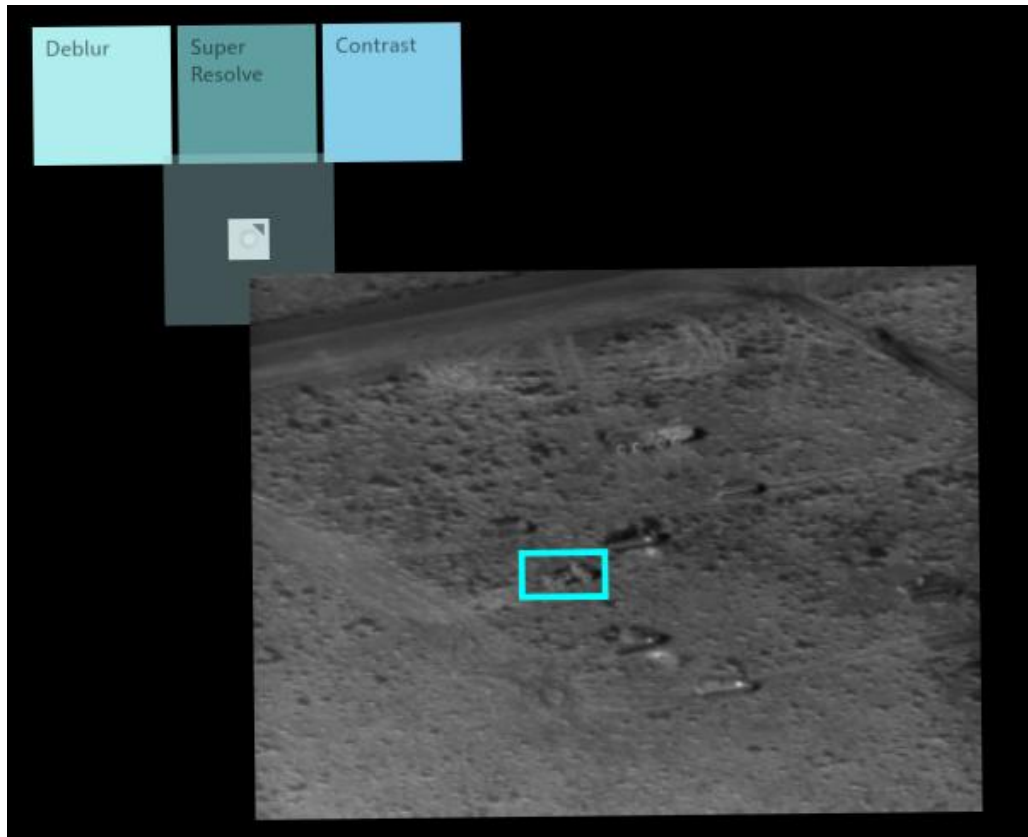
Figure 6. SR Enhancement window.

Once all of the parameters are set appropriately and the region of interest is selected, the SR client application can be executed. To do this, the operator must click or touch on the white square (element menu) in the upper left corner of the Image Enhancement Window then move the mouse or touch device over the SR square and click or touch that square (figure 6).

The results of the SR algorithm will appear in a new SR Image Output window on the main screen in seconds. The length of time to perform the SR is dependent on the number of images, the size of the region to be resolved, the density of the resolution of image and the parameter settings. In addition to the SR Image Output window, a Selected Region window and a Selected Region Gray Scale window will also appear on the main screen. The Selected Region window displays a zoomed in version of the region that was super resolved. The Selected Region Gray Scale window displays a zoomed in version of the region that was super-resolved after converting the image to grayscale. The grayscale image is displayed for comparison since the SR algorithm is applied to grayscale images. These additional windows allow the operator to compare the images, as depicted in figure 7.
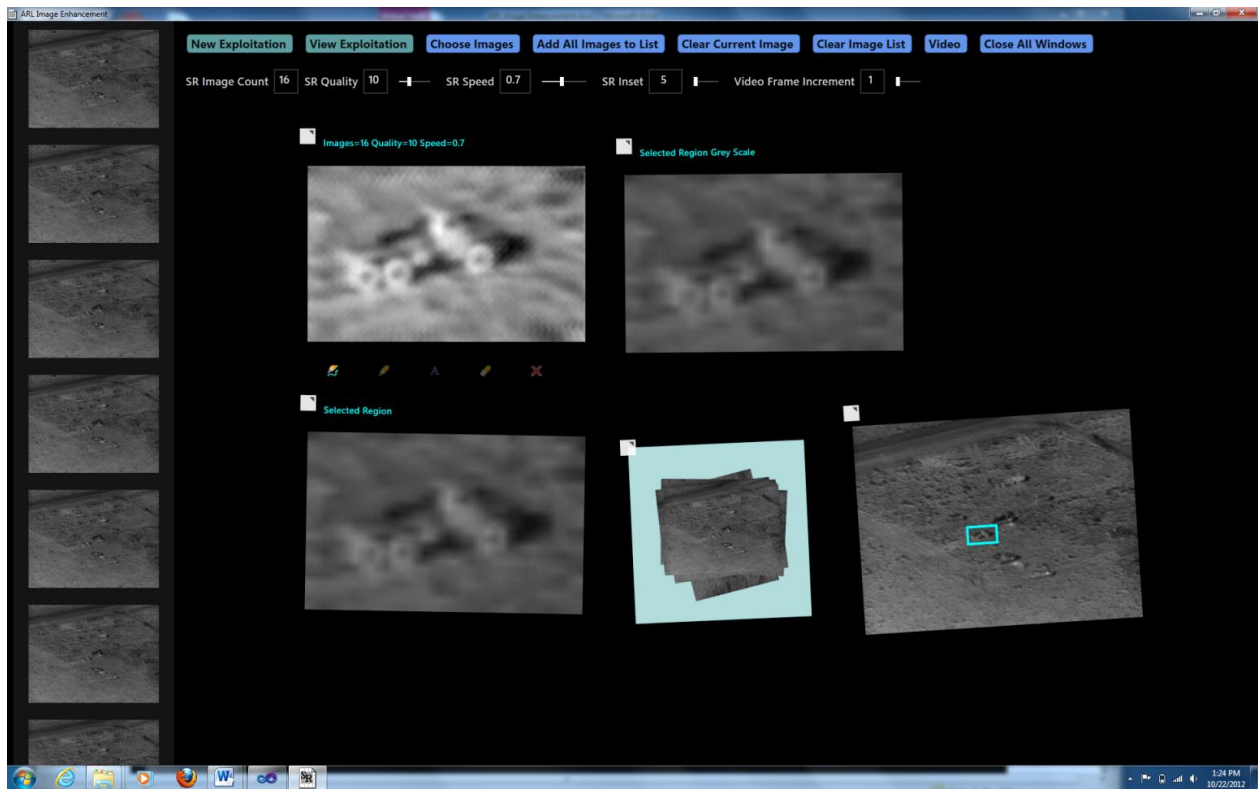
Figure 7. Main screen after SR.

The SR Image Output Window, shown in figure 8, displays the super-resolved image and documents the parameter settings used to achieve the displayed output at the top of the window. The window also allows the operator to annotate the super-resolved image with both drawing and text. When viewing the SR Image Output Window in figure 8, the operator sees a pencil with a blue scribble icon in the lower left corner. The window is now in "move" mode and can be moved around the main screen with the mouse or touch device. If using a touch surface, the window can also be resized as well using the finger pinch. All the windows that appear on the main screen have this move and resize capability. If the touch capability is not an option, the windows can be moved using the mouse but cannot be resized. When the operator clicks with the mouse or touches the pencil with a blue scribble icon, the icon changes to a cross with arrows and the additional four icons are enabled. The pencil icon allows the operator to draw on the image. The A icon allows text to be written on the image where the mouse is clicked or the touch device is touched. The eraser icon erases the drawings that are touched or clicked on with the mouse. To remove the text, the operator must click or touch the text and backspace each letter. The red X icon removes all annotations from the window including text.

Figure 8. The SR Image Window in annotation mode.

3.2 Contrast and Deblur

Two additional algorithms, Contrast (*4*) and Deblur (*5*), can be applied to any image uploaded to the Web service including not only raw images, but also images that have previously been super-resolved, contrast-enhanced, deblurred, and/or annotated. (Applying the Deblur algorithm to a super-resolved image, for example, can further enhance the quality of the image.) The operator can apply these algorithms by clicking or touching the white square or element menu in the upper left corner. Notice that four squares or menu element items appear as shown in figure 9. Selecting the Close element item closes the window. The element item containing the parameter settings for the SR algorithm provides information only and does not execute anything when clicked. Selecting the Deblur and Contrast item elements issues the request to the appropriate Web service specifying the URI of the image currently shown in the window.

Figure 9. Element menu.

The deblur filter is designed assuming the blurring function is Gaussian. Two parameters are used to adjust the deblur filter: mg and alpha. The parameter mg adjusts the magnitude of the inverse Gaussian, and can be varied between 1- and 100-in increments of 1. The parameter alpha adjusts the ratio of the standard deviation of the Gaussian filter to the maximum polar frequency. It controls the bandwidth of the Gaussian filter. The alpha parameter can be varied between 0 and 1 in increments of 0.001. For an initial (default) setting, use alpha=0.25 and mg=5. Decreasing alpha sharpens the image further; however, artifacts may appear with low values of alpha. The operator can tune alpha until the desired sharpness is achieved with minimal presence of artifacts.

When the Deblur element menu item is selected the Deblur Web service is called specifying the URI of current image in the window using the default settings for the alpha and MG parameters. The output of the Deblur Web service appears in a new Deblur window on the main screen, as shown in figure 10. The parameter settings used for the previously applied algorithms are documented at the top of the window. Sliders on the bottom and right side of the output image allow the operator to modify the above stated parameters for the Deblur algorithm. The Deblur Web service executes in real time as the parameters MG or alpha change and a new result image will replace the current result image in the Deblur window. The new parameters are applied to the original image, not the current deblurred image. The Deblur result image can be annotated in the same manner as the SR output image described above. Selecting the Default menu element item allows the operator to execute the Deblur client app using the original default parameters, alpha=0.25 and mg=5.

Selecting the Contrast menu element item executes the Contrast client application on the current image shown in the Deblur Window using default parameter settings a described below. Selecting the Close menu element closes the window without saving.
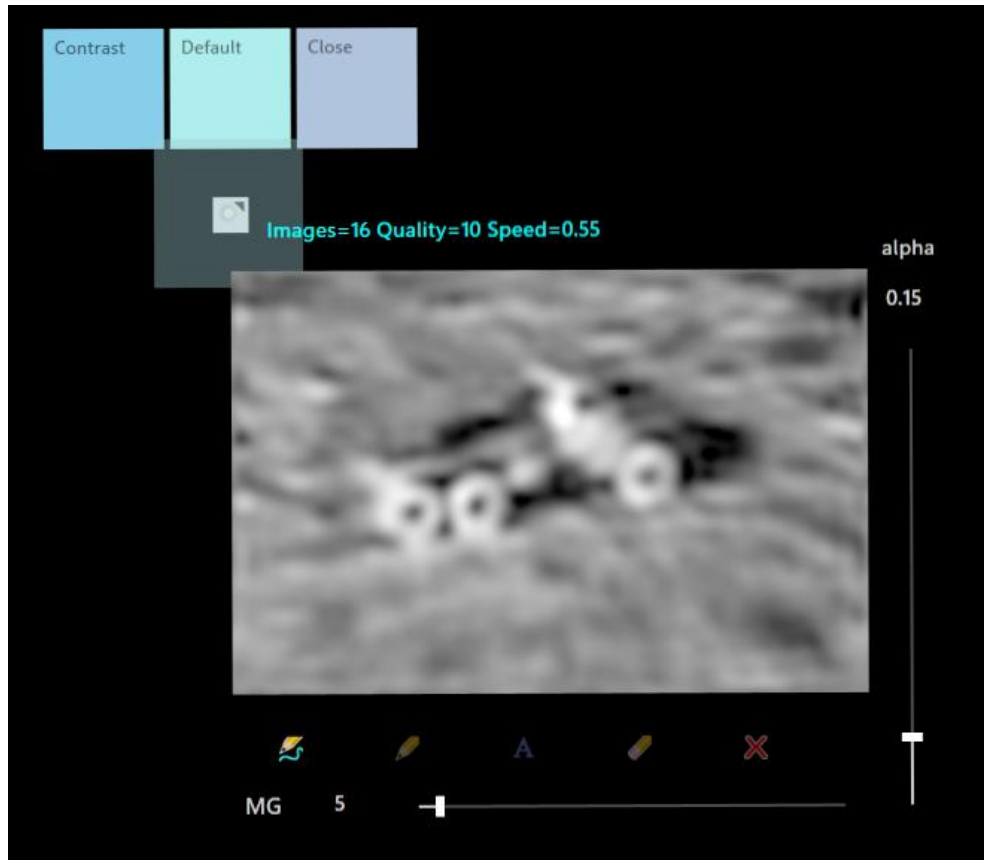
Figure 10. The Deblur window.

The operator also has the option of applying the Contrast algorithm to any image stored on the Web service. The contrast enhancement technique examines the accumulated histogram of pixel intensities in an image. Two parameters, Per Low and Per High, can be adjusted to determine the lower tail and upper tail of the accumulated histogram. The possible range of both of these parameters is between 0 and 1, under the constraint that the sum of per low and per high must always be less than 1. The user interface allows the operator to adjust these parameters as described below.

When the Contrast element menu item is selected the Contrast client application is applied to the current image in the window using the default settings for the Per High and Per Low parameters. The output image of the Contrast client application appears in a new Contrast window on the main screen, as shown in figure 11. The parameter settings used for the previously applied algorithms are documented at the top of the window. Sliders on the bottom and right side of the contrast image allow the operator to modify the above stated parameters for Contrast. The Contrast Web service executes in real time as the operator changes the Per Low and/or Per High parameters and a new image will replace the current image in the Contrast window. The new parameters are applied to the original image not the current contrasted image. This image can also be annotated in the same manner that the SR and deblur output images, as described above.

Selecting the Default menu element item allows the operator to execute the Contrast client app using the original default parameters, Per Low=0.5 and Per High=0.002. Selecting the Dark Image menu element item executes the Contrast application with parameters set to Per Low=0.1 and Per High=0.001. This is should be used when lower pixel values in the image are saturated. Selecting the Light Image menu element item executes the Contrast application with parameters set to Per Low=0.001 and Per High=0.1. This parameter setting is used when higher pixel values in the image are saturated. The Contrast menu element item executes the Contrast client application on the current image showing in the Contrast Window using default parameter settings a described below. Selecting the Close menu element item closes the window without saving.
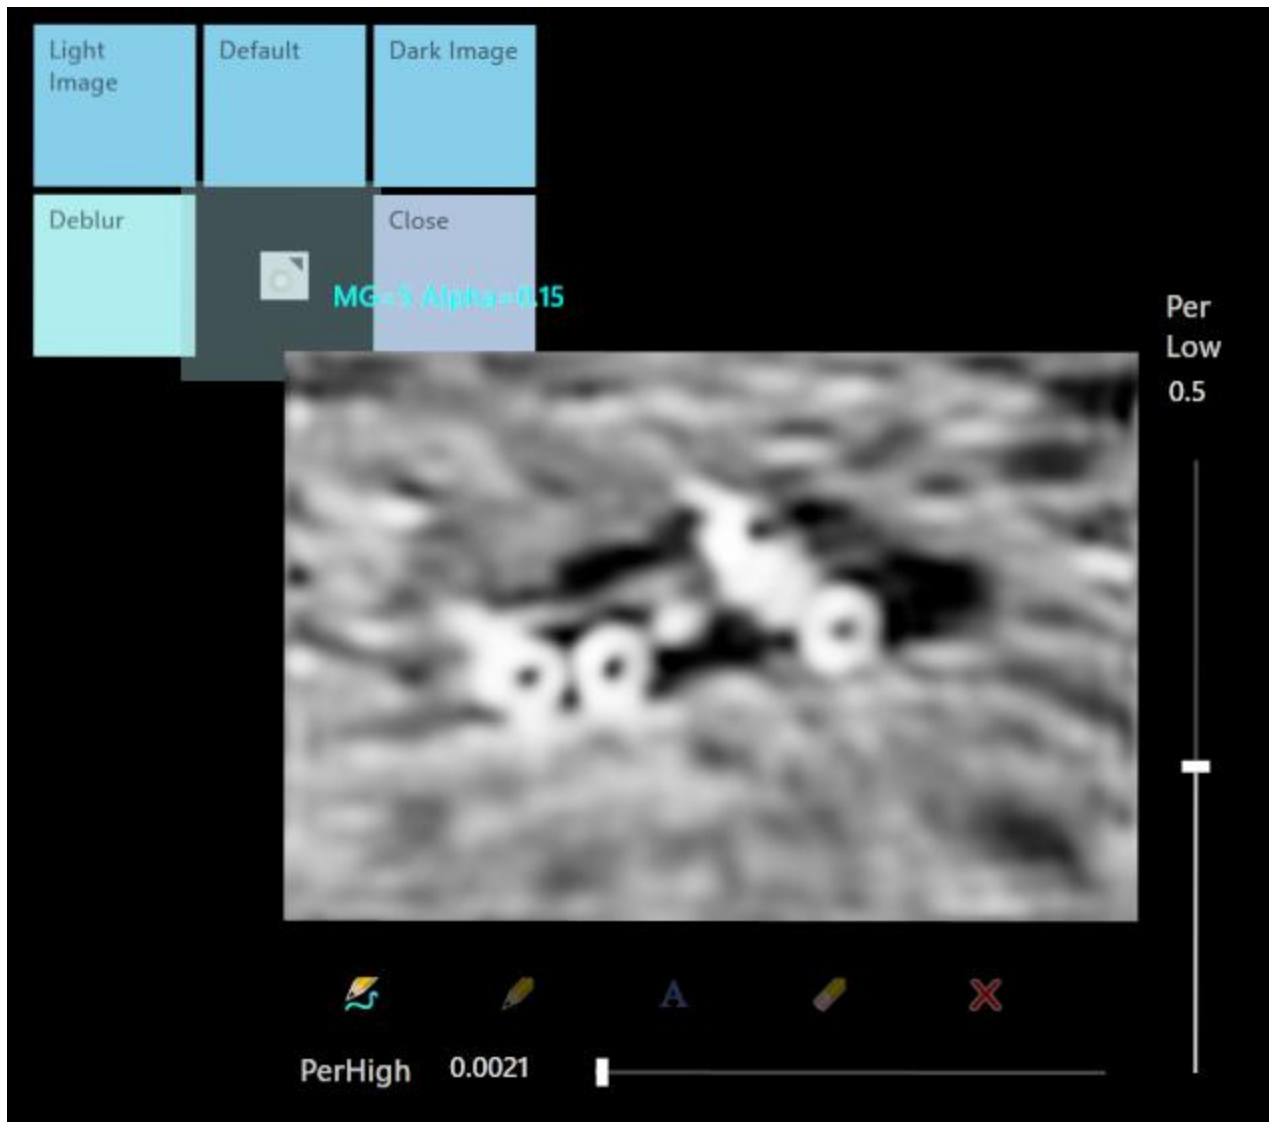
Figure 11. The Contrast window.

3.3 Full Motion Video and Super Resolution

The ARL Image Enhancement Test Bed also allows the operator to view full motion video from various network cameras residing on the roof top of ARL. The operator may also apply the SR application to the full motion video or choose to execute the Force Protection Surveillance System (FPSS) (*5*) tracking algorithm. These operations are described below.

When the blue Video button at the top of the main screen, figure 1, is selected the Select a Video Server pop-up window will appear on the main screen, as shown in figure 12. From this menu, the operator can select from available video servers. The operator can click or touch the server item, which then becomes highlighted.
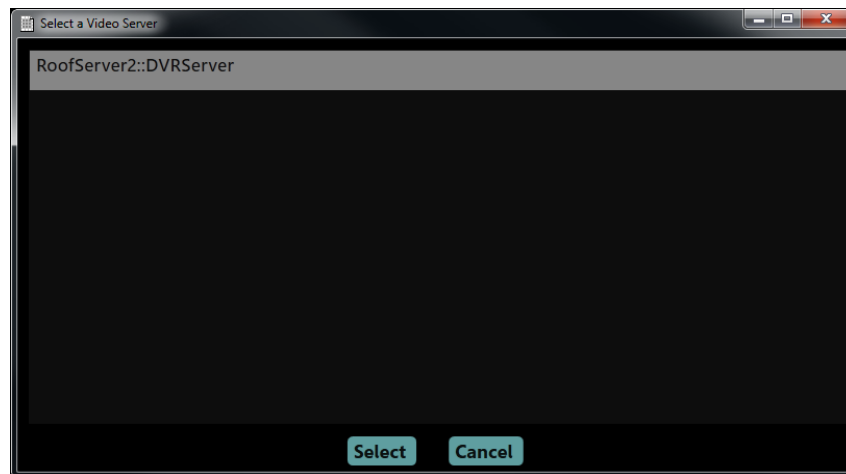
Figure 12. Select a Video Server pop-up window.

When the Select button is activated, the Select a Camera pop-up window appears on the main screen, as shown in figure 13. This displays all the cameras available for the selected server. The operator selects the desired camera and activates the Select button. The full motion video window appears on the main screen, as shown in figure 14. The name of the current camera is displayed in the upper left corner.
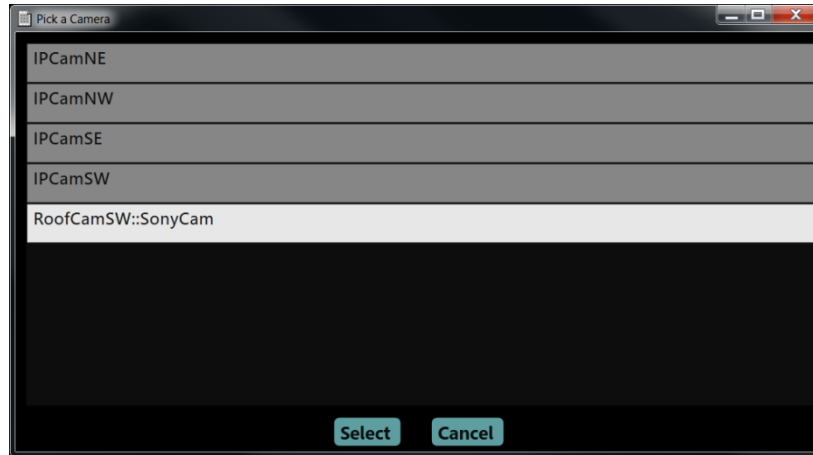
Figure 13. Select a Camera pop-up window.
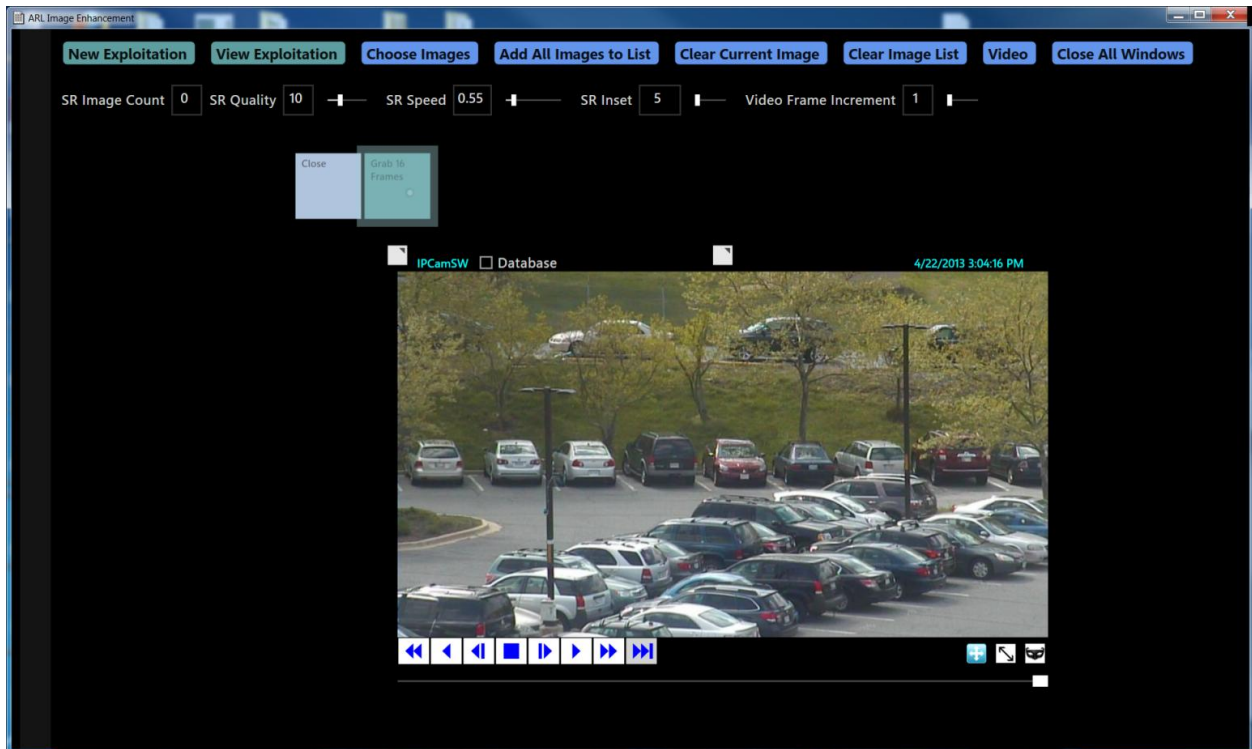
Figure 14. The Full Motion Video window.

The Full Motion Video window contains a variety of different controls. At the bottom of the image, the operator has access to the DVR controls. The DVR stores 48 hours of full motion video. These control buttons are defined as follows starting from the left, fast reverse, reverse, reverse one frame at a time, stop, forward one frame at a time, forward, fast forward, and "live" video. Below these control buttons is a slide bar for controlling where the DVR video is in time. When the video is set to live video mode the slide bar will always be at the far right position.

In the current "move" mode, the operator can move the window within the main screen by clicking and holding the mouse while dragging the window, by touching and dragging the window, or by "flicking" the window with the mouse or touch device. However, when the magnifying glass is toggled to the move arrows, as shown in figure 14, the operator can pan, tilt, and zoom the current camera. These operations can be performed using either the mouse or the touch surface. The wheel of the mouse can be used for optically zooming the camera in and out, changing the field of view. The pan and tilt operations are performed by pressing the left mouse key and dragging it around the video window to the desired viewing region. When using the touch screen the camera is panned by touching the video window and moving the touch device to the desired viewing region. Zooming the camera when using the touch screen is accomplished by using the pinch gesture, pinching together to zoom in and pushing the fingers apart to zoom out.

The button with the arrows icon allows the operator to change the size of the video window. When not using a touch screen, only two sizes are available. The sizes include the default size that originally appears on the screen and an additional size that is equivalent to the pixel dimensions of the actual video image. If a touch screen is used the operator has the ability to make the window smaller and larger by using two or more fingers to pinch and grow the window. The window must be in "move" mode to change the size of the window when using the touch screen unless resizing with the resize button.

In the upper left corner, the element item shows two options, Close and Grab 16 Frames. Selecting Close, closes the video window. Selecting Grab 16 Frames, grabs 16 real-time image frames from the DVR and stores them as a "collection" to the Web service. These images are also displayed on the left side of the main window in the scrolling image list. This allows the operator to execute the SR application on 16 images from the DVR's stored or live full motion video. The operator could also select one frame from the scrolling image list and apply the Contrast or Deblur application to enhance this frame.

The Video Frame Increment slider, shown at the top of figure 14, designates the number of frames to skip between the images being grabbed for SR. This allows the user to potentially add additional movement between the frames if the camera and/or scene is too stable SR to be effective.

3.3.1 Force Protection Surveillance System

The center element menu on the video window (figure 15) allows the operator to select a full motion video tracker. When the FPSS (*6, 7*) tracker element item is selected the FPSS Initialization Parameters popup menu appears on top of the main screen, as shown in figure 16. This menu allows the user to modify the initialization parameters before starting the tracker application. These specific parameters can only be modified prior to starting the FPSS algorithm. The FPSS Initialization Parameters are defined below. To start the FPSS with the default parameters selected, hit the OK button. If the Cancel button is selected, the tracker will not be

started and the menu will disappear. Each of the FPSS Initialization Parameters (*8*) are described in detail below.
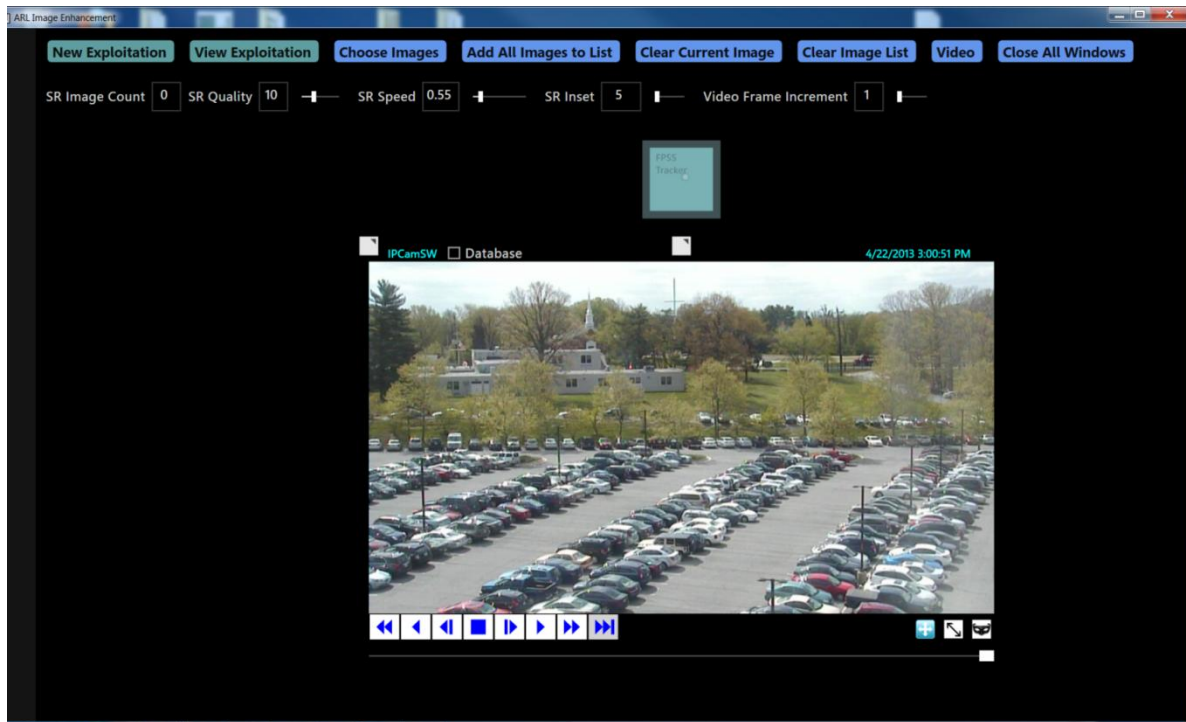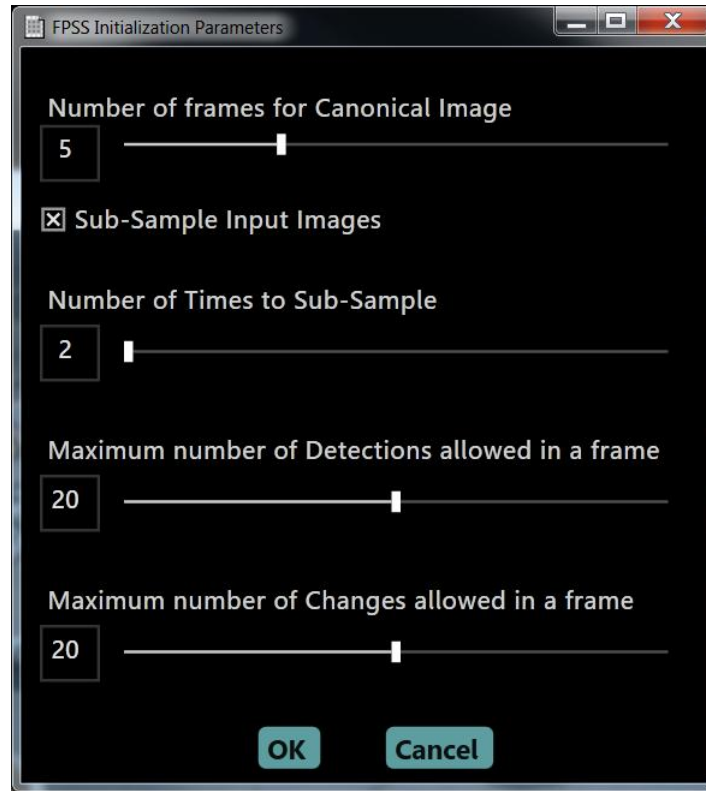

Figure 15. Tracker selection.

Figure 16. FPSS Initialization parameters.

3.3.2 FPSS Initialization Parameters

The Number of Frames for Canonical Image values range from 3 to 10 frames, with a default value of 5 frames. This parameter sets the number of frames used to construct a representative background model for a given interval. A larger value increases the requirements to store and compute more frames but provides a more stable background model. A smaller value is recommended for a faster changing background or a faster rate of input frames from a video stream.

The Sub-Sampling Input Images parameter is a true false value, with the default set to true. This parameter determines whether the input images will be processed at their original size or a smaller size. If the box is checked setting this parameter to true, the size of the input are reduced by the factor specified in the Number of Times to Sub-Sample parameter, before the images are further processed. While saving considerable computational resources, the sub-sampling feature does not introduce significant degradation in detection and tracking accuracy if moving objects are still clearly noticeable in the sub-sampled images. Sub-sampling is recommended.

The Number of Times to Sub-Sample parameter ranges from 2 to 5 in increments of 1, with the default set to 2. The operator would only use a value of 5 in the case of very large input images containing large moving objects. With a value of 2, the input images are reduced to a quarter of

the original size before processing. A value of 5 reduces the size of the input images by 1/25 of the original size.

The Maximum Number of Detections Allowed in a Frame parameter values range from 10 to 30, with a default value of 20. This parameter specifies the maximum number of moving objects allowed in the tracking process at any given time. The operator should be aware that the higher the number of tracked objects, the more computational resources will be required.

The Maximum Number of Changes Allowed in the Frame parameter values range from 10 to 30, with a default value of 20. This parameter specifies the maximum number of slow changes allowed in a given scene. It is the counter part to the Maximum Number of Detections Allowed in the Frame above. As with the previously described parameter, the operator should be aware that the higher the number of slow changes, the more computational resources will be required.

Clicking the Cancel button or the red X in the upper right-hand corner will close the FPSS Initialization Parameters Window without starting the FPSS Tracker. If the OK button is selected, the tracker algorithm will begin processing the images.

The tracked objects in the full motion video window are denoted by colored track boxes with associated ID numbers, as shown in figure 17. The tracker also performs a limited amount of object classification, which is designated by the color of the boxes surrounding the object and its associated ID numbers. Pedestrians are designated with the color cyan, vehicles are designated with the color blue, and unknown moving objects are designated with the green.

Once the tracker process begins, an additional menu element appears on the right upper corner of the video window. The element menu displays five additional options for the FPSS when clicked, as shown in figure 17. These options are defined below.
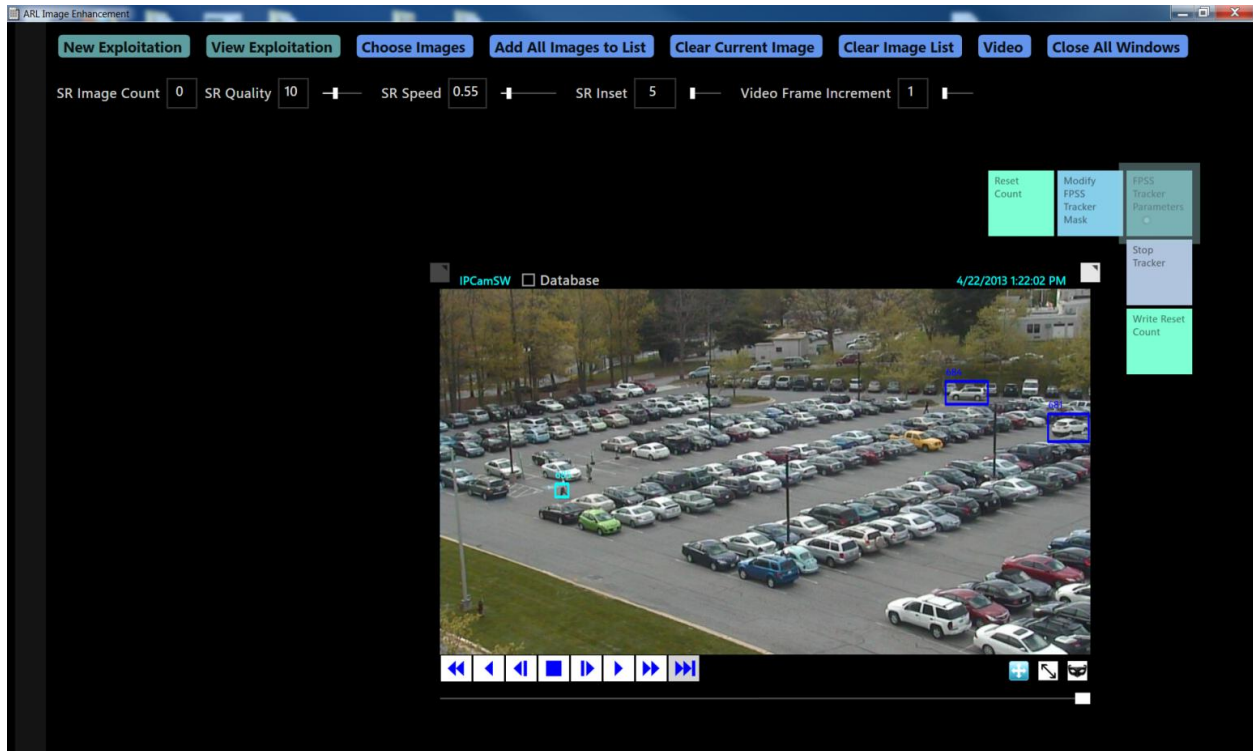
Figure 17. Video window with FPSS tracker running and additional tracker options.

When the option Modify FPSS Tracker Mask is selected, the pop-up menu shown in figure 18 appears on top of the main window. This menu allows the operator to develop a mask for the tracker to achieve better performance and fewer false alarms. When the Add Critical Region/Regions for detection is selected, the operator can draw on the video window in red designating areas in which the tracker parameters will be enhanced, creating a critical region. The alert level of this region is increased by a Critical Boost Factor multiplier (figure 21). In the example shown in figure 19, the critical region allows tracker to be more sensitive in an area that is farther away and the moving objects are smaller. When the Add Region/Regions to be Ignored is selected, the operator can draw on the video window in yellow, designating areas in which the tracker will ignore all motion in this area thus removing all detections. When two regions overlap, in which one is a critical region and the other is an ignore region, the top-most region or the region most recently added takes precedence in the area of overlap. When the Erase Regions is selected the critical or ignore region that is selected will be erased and set back to the default parameters. The Turn Off Modify Map Mode allows the operator to interact with the window as usual. Selecting Close, sets the Turn Off Modify Map Mode parameter to true and closes the pop-up menu leaving the tracking mask as it has been designed. The Reset Count and Write Reset Count options pertain to the Database capability when the Database checkbox is checked and is described later.
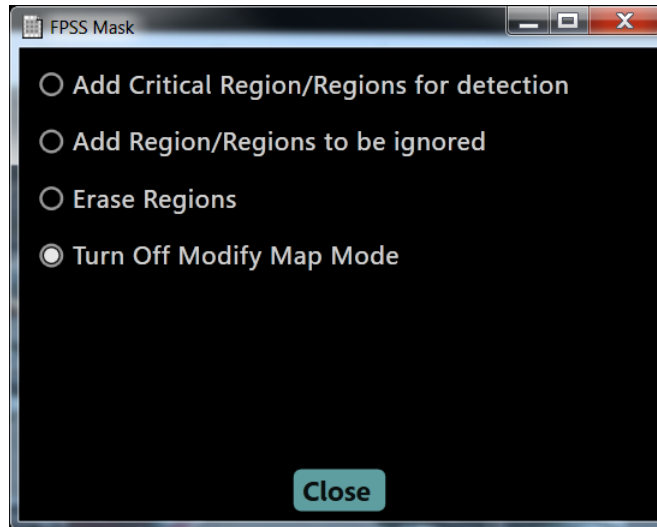
Figure 18. FPSS Mask menu,

The Red and Yellow tracker mask regions are shown in figure 19. Once the tracking mask is complete, the operator can hide/show the mask by selecting the mask toggle button at the bottom right of the Video Window, also shown in figure 19.
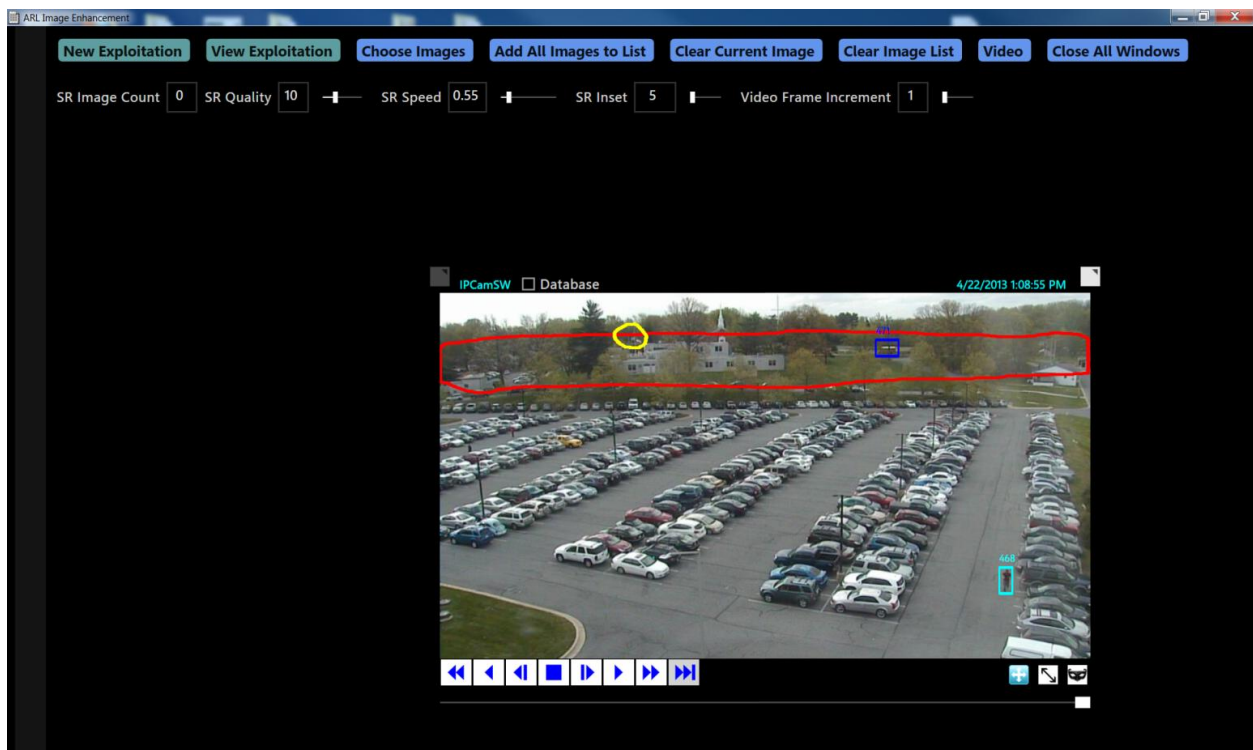

Figure 19. FPSS tracker mask.

19

3.3.3 FPSS Parameters

When the FPSS Parameters (*8*) Menu Element Item is selected, the parameters specific to the FPSS can be modified in real time while the tracker is running. The pop-up menu appears on top of the Main window as a tabulated window. The parameters for each tab are described in the subsections below. Note that the sliders allow the parameter to be modified by moving the slider or typing into the corresponding text box.

3.3.3.1 The Image Tab

The Show Slow Changes in the Scene parameter is a true/false value, with a default of true. When the check box is checked setting the parameter to true, the slow changes in the input scene are shown. These changes are based on the differences between the current scene and the scene acquired when the background was last manually reset by pressing the Update Background button on the Background Tab.

The Median Filter the Input Image parameter is a true/false value, with a default of false. When the checkbox is checked or unchecked, the median filtering of the input image to reduce noise is enabled or disabled. The median filter is used to remove the "salt and pepper" noise that can be introduced by a faulty camera or noisy communication link. The filter is disabled by default since it uses a sizeable amount of computation time.

The Pixel Floor the Input Image parameter is a true/false value, with a default of false. The parameter is used to control the brightness of the image. When the checkbox is checked, setting the parameter value to true, the minimum pixel value of each input image is shifted to zero. This feature is helpful when the input video streams have significantly different dynamic ranges and minimum pixel values. It allows the subsequent computations and comparisons to be made on an equal footing. This problem is more common on a mixed-mode network, in which color and forward-looking infrared (FLIR) cameras co-exist in the same sensor network.

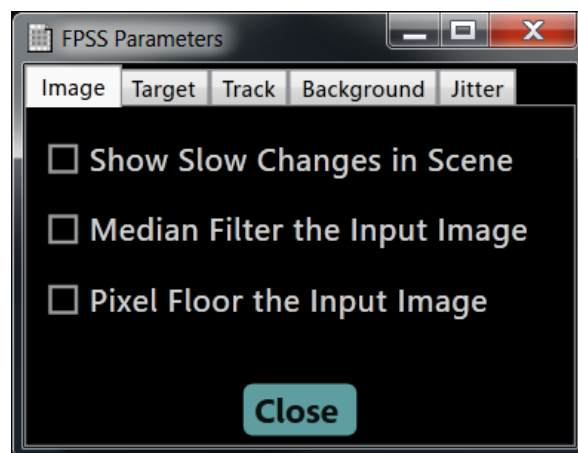Figure 20 shows the image parameters tab.

Figure 20. Image parameters.

3.3.3.2 The Target Tab

The Acceptance Threshold of Motions parameter values range from 500 to 3000 with a default value of 2000. This parameter is a threshold value used to declare a motion at a given location on the input image. If the activation value computed at that pixel location is higher than this threshold value, a motion is detected. Lower threshold values increase detections of motion but also increase false alarms. The allowable parameter values are those typical for detecting automobiles and people moving on a nearby parking lot. This threshold must be larger than the Object Edge Threshold of Motions parameter described above in order to have a non-zero object size. This constraint is enforced by the sliders in the parameter menu.

The Object Edge Threshold of Motions parameter values range from 100 to 1000 with a default value of 100. This parameter is used to delineate the boundary of a moving object based on the activation value of the boundary pixels. This threshold must be smaller than the Acceptance Threshold of Motions parameter described above in order to have a non-zero object size. This constraint is enforced by the slider in the parameter menu. The greater the difference between the Acceptance Threshold of Motions parameter value and the Object Edge Threshold of Motions parameter value, the larger the bounding box that surrounds the moving object will become.

The Critical Region Boost Factor parameter values range from 2.0 to 10.0 with a default value of 3. This parameter sets the multiplier for the critical regions defined by the mask. Implementation of the critical region mask is described above. The activation value within the critical regions will be multiplied by the boost factor with a value of greater than one. Therefore, less significant motions within the critical regions will have a greater chance to meet the motion acceptance threshold and be declared as a valid moving object.
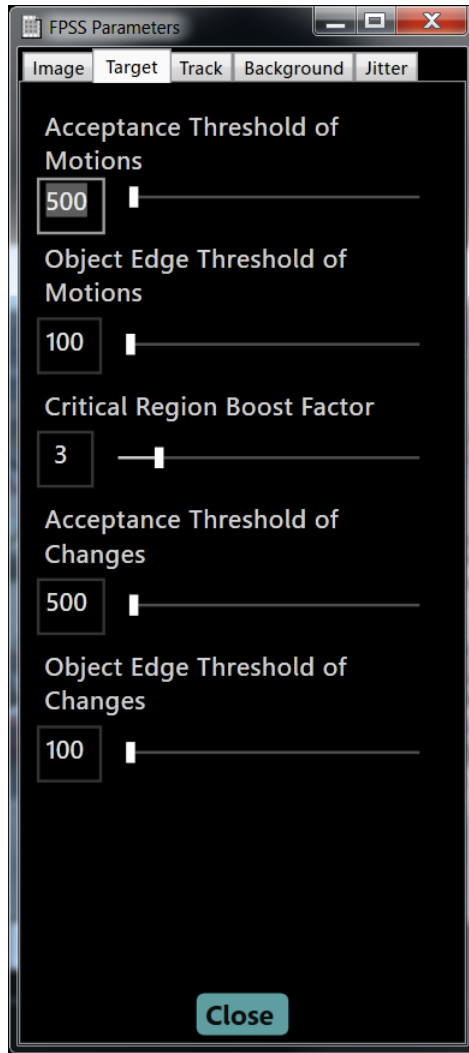
Figure 21 shows the Target tab.

Figure 21. Target parameters.

The Acceptance Threshold of Changes parameter values range from 500 to 3000 with a default value of 500. It is the counterpart of the Acceptance Threshold of Motions for detecting slow changes in the input scene. This threshold parameter is used for detecting slow changes in the input scene. The value of this threshold also depends on the size and contrast of the scene changes anticipated for a given scenario. This threshold must be larger than the Object Edge Threshold of Changes parameter described above in order to have a non-zero object size and is enforced by the slider in the parameter menu. This constraint is enforced by the slider in the parameter menu.

The Object Edge Threshold of Changes parameter values range from 100 to 1000 with a default value of 100. It is the counterpart of the Object Edge Threshold of Changes for defining the size of slow changes. This threshold must be smaller than the Acceptance Threshold of Changes parameter described above and is enforced by the slider in the parameter menu. This constraint is enforced by the slider in the parameter menu.

3.3.3.3 The Track Tab

The Matching Bound for Proximity values range from 0.5 to 5 in increments of 0.5 with a default value of 2. This parameter is used to set up the search boundary for a given moving object in the next input frame based on the current location and target size. For example, when the parameter is set to a value of 2, a moving object is expected to move in the subsequent frame within a region that is centered at the current location and having twice the current target dimensions. As this value increases, a larger search area for the next location of a moving object is generated. A large value for this parameter will reduce the possibility of losing track of an object but will also increase computation time and the risk of mismatching a moving object with other nearby moving objects. Figure 22 shows the Track tab.

Figure 22. Track parameters.

The Min Scope of Proximity values range from 10 to 50 with a default value of 20. This parameter sets the minimum size, in terms of pixels in each dimension, of the search space for a

moving object in the subsequent frame. This parameter compensates the Matching Bound for Proximity parameter, so that the search boundary will not be set to an unrealistically small size for smaller moving objects. Otherwise, small but rapidly moving objects will not be tracked consistently.

The Min Target Size parameter values range from 10 to 100 with a default value of 20. This parameter determines the minimum size of a legitimate moving target. The target size is computed by multiplying its estimated height and width. Any object smaller than this size is ignored by the tracker.

The Min Average Velocity parameter values range from 0.5 to 3 in increments of 0.5 with a default value of 0.5 pixels per frame. This parameter sets the threshold for the minimum average velocity of a transient object in order to be considered for further tracking. The minimum average velocity is defined in terms of the average number of pixels per frame moved in a given number of frames. The value for the number of frames is determined by the Number of Frames for an Average Canonical Image parameter (default = 5) on the FPSS Initialization Parameters window, which can only be modified prior to starting the tracker.

The Min Frames for Transient Detections parameter values range from 1 to 5 frames with a default value of 3 frames. This parameter sets the threshold for the minimum number of frames in which a transient moving target is consistently detected before this transient target is officially recognized as a legitimate moving target and further tracked in future frames. This is used to help exclude short-lived transient objects or noise in the tracking process, reducing the number of false detections.

The Max Frames Allowed for Stopped Detections parameter values range from 5 to 20 with a default value of 10 frames. This parameter sets the threshold for the minimum number of non-detection frames required to designate a previously moving target as a stopped object. Once a target is declared stopped, it is no longer actively tracked and reported by the tracker. The tracking history of this target is still maintained and can be continued once it begins moving again.

The Min Frames before Erasing Det parameter values range from 10 to 50 with a default value of 30 frames. This parameter sets the threshold for the number of frames required before erasing a stopped object from the tracking history. Once the object is erased from the tracking history, it will be considered a new moving object by the tracker when it begins moving again.

3.3.3.4 The Background Tab

The Min Pixel Value Difference for Detection parameter values range from 10 to 50 with a default value of 20. This parameter sets the threshold for the minimum average pixel value difference between a region of interest and its background pixels. A lower value of this parameter will enable more regions to be examined for potential detections of movement or

change. This parameter is used in conjunction with the Acceptance Threshold of Motions and the Min Frames Required for Transient Detections parameters to declare a new or ongoing detection.

The Min Pixel Stability Index Required parameter values range from 2 to 5 with a default value of 2 frames. This parameter sets the threshold for the minimum number of frames required for a stable pixel to be considered as part of the background. It is used to separate background pixels from foreground pixels. Its value should not be set too high, especially when the background is changing from time to time, as when the camera is operating in step-stare mode.

The Max Pixel Stability Index Allowed parameter values range from 30 to 100 with a default value of 30 frames. This parameter determines the maximum number of frames that a stable pixel can keep in its memory. A moderate value for this parameter will provide a consistent background model that can withstand occasional disturbances in pixel values while maintaining the flexibility to reassign a background pixel to a foreground pixel when persistent and substantial fluctuations in the pixel occur.

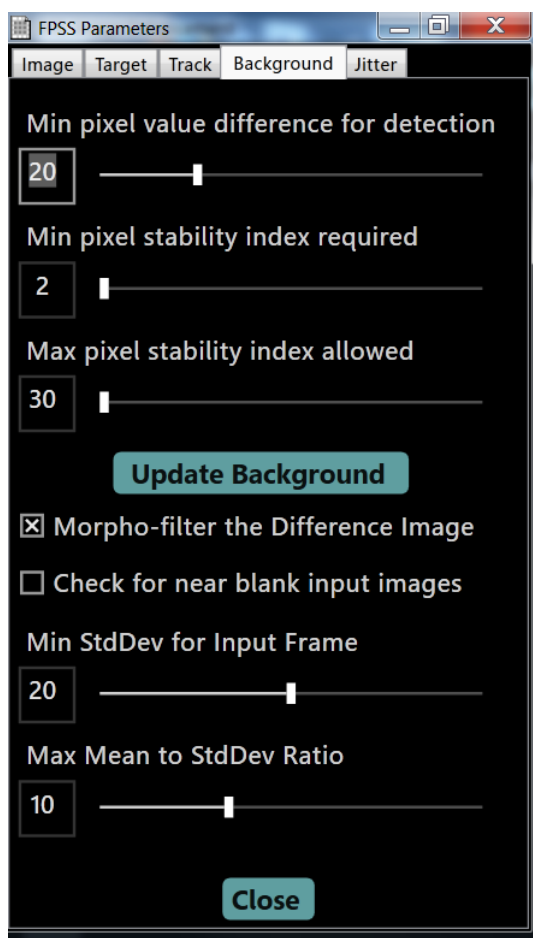Figure 23 shows the Background tab.

Figure 23. Background parameters.

When the Update the Background button is clicked, the background model for change detection will be replaced with the current scene in the current image frame. When the background is updated in this way, all the change detections observed between the previous background and current background scene will be erased and the change detection process is restarted. This will occur each time the Update Background for Changes button is clicked.

The Morpho-Filter the Image parameter is a true-false value, with the default of true or checked as shown in the checkbox. It is used to determine whether or not a morphological filter will be included as part of the tracking process. This filter is used to remove noise and smooth the difference images. This filter should only be included when the scene contains noisy or fragmented targets, since it requires additional computation.

The Check for Black Images parameter is a true-false value, with a default of false or unchecked as shown in the checkbox. It is used to check if the input images are blank in nature, such as those produced by dead cameras or faulty transmission lines.

The Min StdDev for Input Frame parameter values range from 10 to 30 with a default value of 20. This parameter is used to determine whether or not an input image has sufficient variability in its content, so as to justify further processing. If the standard deviation of the input image is below the minimum standard deviation threshold, the image will be skipped.

The Max Mean to StdDev Ratio parameter values range from 5 to 20 with a default value of 10. This parameter is used to detect and avoid problematic input images by checking the ratio between the mean and standard deviation of each input image. A small mean with a very large standard deviation indicates a large fluctuation of pixel values around zero. Whereas, a very large mean with a small standard deviation may indicate bright pixels with relatively small variation. In either case, these apparently abnormal input images will not be considered in further processing.

3.3.3.5 The Jitter Tab

The Check for Jitter on Input Images parameter is a true-false value, with the default set to false or unchecked as shown in the checkbox. This parameter is used to enable/disable the jitter checking process for the input images. Notice in figure 24 that when the Checkbox for this parameter is unchecked, the remaining parameter sliders are disabled since the jitter checking process will not be performed. When this Checkbox is checked the remaining parameter sliders become enabled allowing the operator to adjust the additional parameters for the jitter checking process. The jitter checking process registers the input images against previous frames to generate a stable background in order to provide meaningful foreground information. This process is only recommended when there is pronounced jitter of the camera system since the image registration process takes a substantial amount of computational resources.
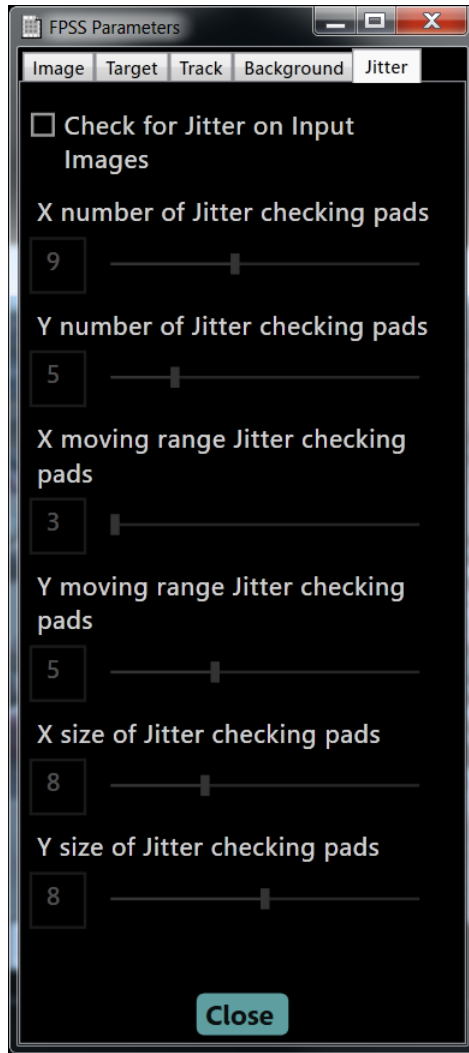
Figure 24. Jitter filter parameters.

The additional jitter checking parameters are described below. They can be modified using the slider or corresponding text box.

The X number of Jitter checking pads parameter values range from 5 to 15 in increments of 2, since an odd number is required, with a default value of 9. This parameter sets the number of jitter sensing and estimation pads or probes in the horizontal direction. It is dependent on the values of "X moving range of jitter checking pads" and "X size of jitter checking pads" described below. When setting this parameter, these values should also be considered in order to set the number of horizontal jitter checking pads to be large enough to cover a representative region of the input image in the horizontal direction. However, this region should not be so large that the checking process cannot be accomplished efficiently and robustly. An odd number is required for this parameter value to ensure that a pad exists near the vertical center line of the input image.

The Y number of Jitter checking pads parameter values range from 3 to 13 in increments of 2, since an odd number is required, with a default value of 5. This parameter is the counterpart of "X number of jitter checking pads" in the vertical direction and subjected to the effect of "Y moving range of Jitter checking pads" and "Y size of Jitter checking pads" below. An odd number is required for this parameter value to ensure that a pad exists near the horizontal center line of the input image. These parameter values are smaller than the X number of jitter checking pads since typical images have a smaller height than width measurement.

The X moving range of Jitter checking pads parameter values range from 3 to 9 in increments of 2, since an odd number is required, with a default value of 5. This parameter is used to determine the range of pad movement in the horizontal direction. A larger range of movement or parameter value is required to correct for a scenario containing a large amount of jitter at the expense of a higher computational cost. This additional computational cost of having a very large range of pad movement can be offset by decreasing the number and size of the jitter checking pads.

The Y moving range of Jitter checking pads parameter values range from 3 to 9 in increments of 2, since an odd number is required, with a default value of 5. This parameter is the vertical counter part of the X moving range of jitter checking pads parameter described above. It is used to determine the range of pad movement in the vertical direction.

The X size of Jitter checking pads parameter values range from 5 to 20 pixels in increments of 1 with a default value of 8. This parameter sets the horizontal size of the jitter pad or probe in terms of the number of pixels. The value should be large enough for a meaningful comparison between image patches from two different frames, but small enough to achieve a robust and efficient matching process. The value is dictated by the representative background features or objects of given scenario. Scenes in which larger objects are being tracked require larger pad sizes or values of this parameter.

The Y size of Jitter checking pads parameter values range from 5 to 20 pixels in increments of 1 with a default value of 8. This parameter is the vertical counter part of the X size of jitter checking pads parameter described above.

3.4   View Exploitation

Looking back to the ARL Image Enhancement initial screen shown in figure 1, three additional buttons at the top of the main screen need to be described. First, the blue Close All Windows button closes all open image windows on the main screen with the exception of the full motion video windows. When this button is used to close the image windows on the main window, the results of the image processing algorithms and corresponding annotations are not saved to the Web service. However, if the operator selects the green New Exploitation button, the previously mentioned input image collections, their image processing algorithm results along with various annotations will be saved to the Web service for later analysis before closing the windows. The

green View Exploitation button allows the operator to view previously saved exploitations. When this button is selected the window shown in figure 25 appears on top of the main window.
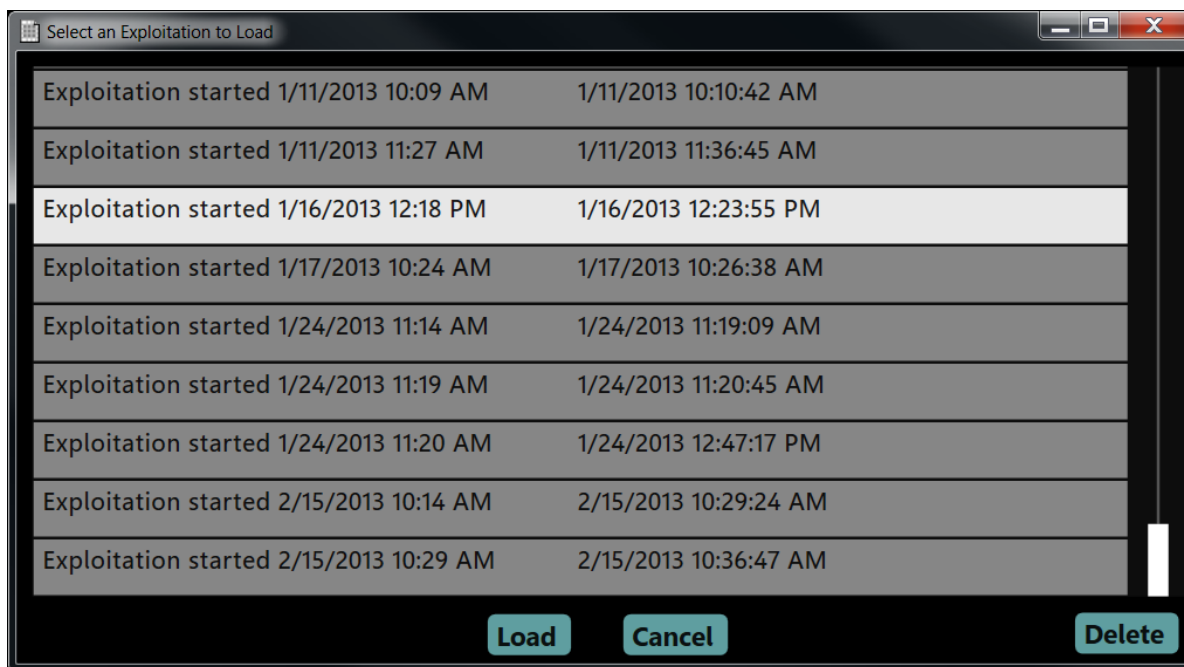

Figure 25. Select an Exploitation to Load window.

In this window, the operator can scroll through a list of Exploitation Collections residing on the Web service. The operator highlights the exploitation of interest with the mouse or touch device. The operator has the option to select Cancel, closing the Select Exploitation to Load window and returning to the main ARL Image Enhancement window. The operator also has the option to select the Delete button to permanently remove the selected exploitation collection from the Web service. This is used for Web service maintenance. When the operator selects the Load button, the View Exploitation window shown in figure 26 appears on top of the main ARL Image Enhancement window. This window can be resized or moved allowing the main ARL Image Enhancement window to be viewed simultaneously.
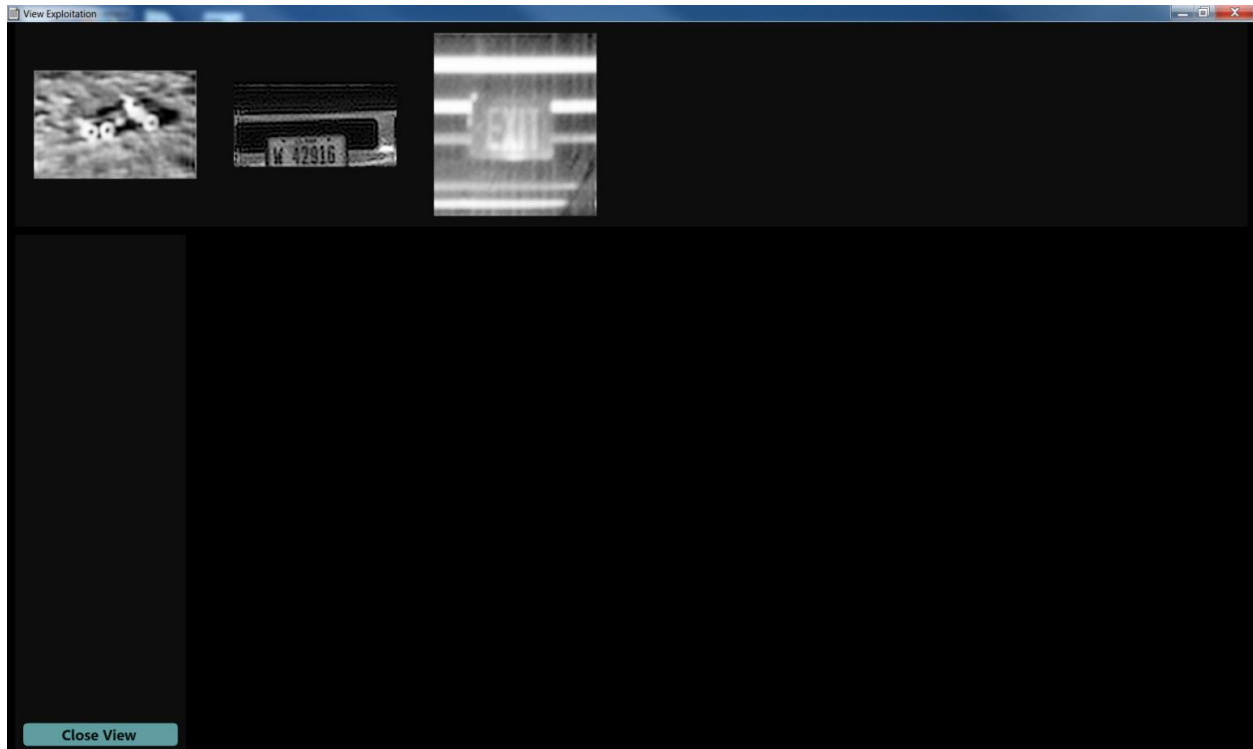
Figure 26. Initial View Exploitation window.

The images in the horizontal scroll bar across the top of the Select Exploitation window shown in figure 26 are the three image collections contained in this specific exploitation. When the operator selects the first image, a list of the image results from each of the various image processing algorithms executed on the original image or images appears in the vertical scroll bar on the left side of the window, as shown in figure 27. Each result image in the list is labeled with the name of the corresponding image processing algorithm that had been applied to the input image/images.
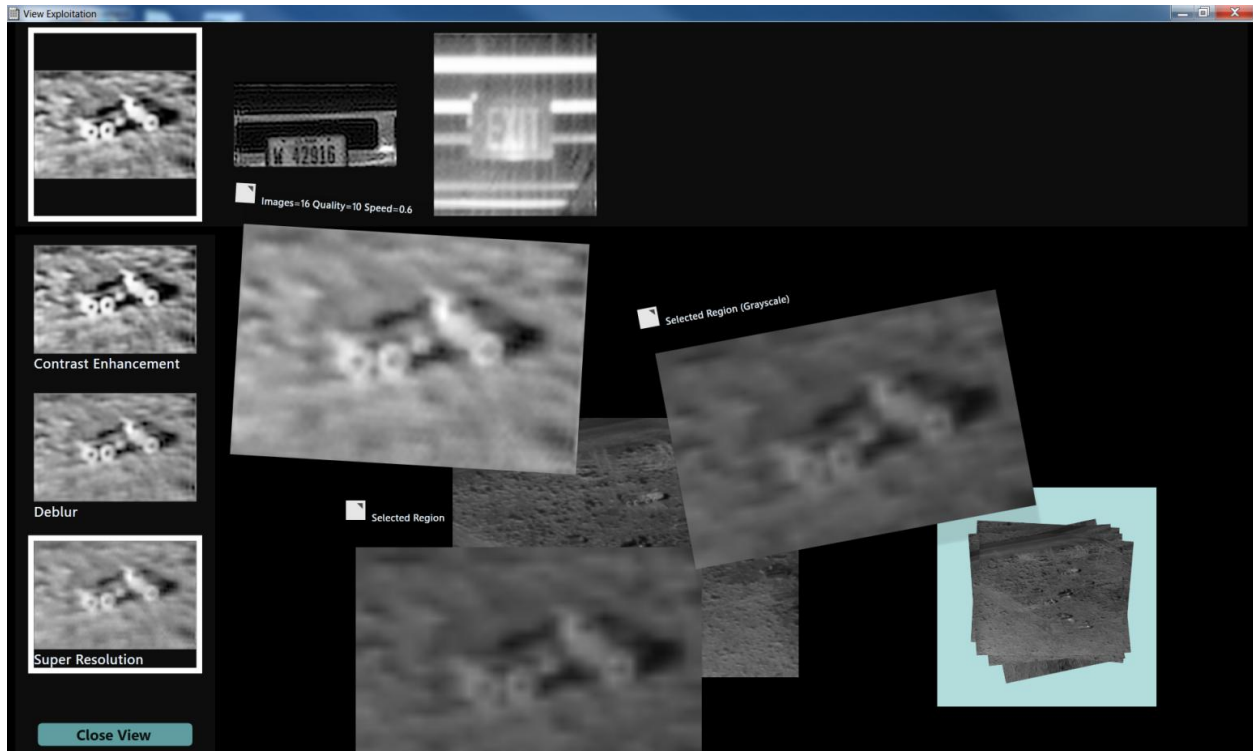
Figure 27. View Exploitation window showing the SR results collection.

In the example shown in figure 27, the image collection labeled Super Resolution has been selected. Selecting this image opened five new windows in the center of the View Exploitation window. These windows contain all the images and windows that pertain to SR, including the original image with the blue rectangle box surrounding the selected region to be super-resolved, the stack box with the blue background containing all the images used to perform SR, the selected region, the selected region in grayscale, and the results of the SR algorithm. Notice that the result of the SR algorithm is labeled with the settings for the parameters used for SR. These windows are for analyst viewing only and additional image processing cannot be performed on these images nor can additional annotations be added. Each window can be closed using the menu element Close button shown in figure 28. Figure 28 also shows the results of when a Contrast Enhancement collection is selected.
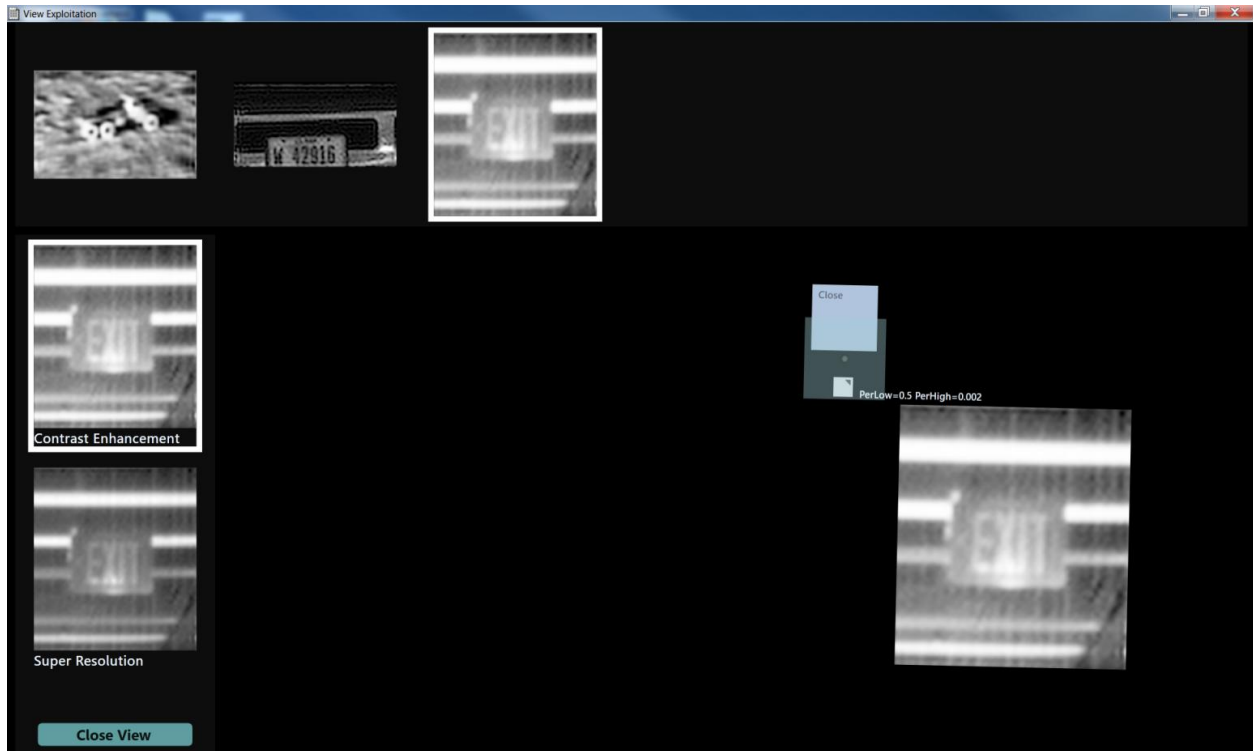
Figure 28. View of the Exploitation window showing the Contrast Enhancement results.

Figure 29 gives an example of viewing two image collections within the exploitation simultaneously. The operator does not need to close all the windows opened while viewing one image collection before opening an additional image collection selected from the horizontal scroll bar.

Figure 29. View of the Exploitation window showing results from multiple image collections.

The operator can view multiple View Exploitation windows by going back to the ARL Image Enhancement window and selecting the View Exploitation button addition times. The Close View button closes the entire Exploitation window. This window can also be closed by selecting the red x in the upper right corner.

3.5 Exiting the ARL Image Enhancement Application

Simply click or touch the red x in the upper right corner of the main ARL Image Enhancement window to exit the ARL Image Enhancement application. Upon exit the existing Super Resolution, Contrast, and Deblur windows will be closed without saving their results to the Web service. All open video windows with or without trackers in operation will also be closed upon exit, gracefully stopping the trackers and closing any additional menu pop-up windows associated with the trackers if necessary.

4. Recent Modifications

The ARL Image Enhancement Test Bed was recently modified to support data collection efforts by ARL's Sensor and Electronic Devices Directorate (SEDD) for other research purposes (anomaly detection and classification, for example). The ARL Image Enhancement Test Bed integration software was modified to provide ground-truth data for this effort by keeping track of

the number of vehicles and pedestrians detected in each direction on a road using output information from the FPSS tracker. The total number of detections in each direction, which includes vehicles and pedestrians and unknowns, is also calculated. A database was designed and developed using MS SQL Server to store the required data. The data are written to the database automatically every 5 min. These data include the number of pedestrians detected in each direction; the number of vehicles detected in each direction; the total number of detections in each direction; the timestamp of the first image frame for the 5-min period; and the timestamp of the last image frame for the 5-min period. The video frames are also recorded in the DVR for ground-truth purposes. Storing the timestamps allows the researcher to easily correlate the areas of interest between the various sensors, the database, and video.

5. Conclusion

The ARL Image Enhancement Application is still under development. It is designed to be extensible to allow for the addition of new image processing algorithms as they become available. Other full motion trackers can also be added to the video window options menu and executed in a similar manner to the FPSS with its own unique parameters in a pop-up window. It is currently under consideration to record the motion detections or events from the FPSS in a Terra Harvest xml schema format. These events could then be stored in a database or disseminated to network consumers in real time. The database could be used for analysis by anomaly detection applications as well as any other analyst applications.

6. References

1. Winkler, Robert P.; Schlesiger, Chris. *Image Processing REST Web Services*; ARL-TR-6393; U.S. Army Research Laboratory: Adelphi, MD, March 2013.

2. Young, S. Susan; Driggers, Ronald G. Superresolution Image Reconstruction From a Sequence of Aliased Imagery. *Applied Optics* **2006**, *45* (21), 5073–5085.

3, Driggers, Ronald G.; Krapels, Keith A.; Murrill, Steven.; Young, S. Susan; Theilke, Matthew.; Schuler, Jonathan M. Superresolution Performance for Undersampled Imagers. *Optical Engineering* **2005**, *44* (01).

4. Young, S. Susan.; Driggers, Ronald G. Signal Processing and Performance Analysis for Imaging Systems. *Artech House*, 2008.

5. Young, S. Susan.; Driggers, Ronald G.; Teaney, Brian P.; Jacobs, Eddie L. Adaptive Deblurring of Noisy Images. *Applied Optics* **2007**, *46* (5), 744–752.

6. Chan, Alex Lipchen. *Force Protection Surveillance System: Algorithm and Performance*; ARL-TR-5322; U.S. Army Research Laboratory: Adelphi, MD, 2010.

7. Chan, Alex Lipchen. A Robust Target Tracking Algorithm for FLIR Imager. *Proceedings of the SPIE*, Volume 7696, pp. 769603-769603-11 (2010).

8. Chan, Alex Lipchen. *ARL FPSS Parameter Descriptions*. Unpublished manuscript, 2008.

List of Symbols, Abbreviations, and Acronyms

ARL   U.S. Army Research Laboratory

CPU   central processing unit

FLIR   forward-looking infrared

FPSS   Force Protection Surveillance System

GUI   graphical user interface

REST   representational state transfer

PC   personal computer

SEDD   Sensors and Electronic Devices Directorate

SR   Super Resolution

WPF   Windows Presentation Foundation

INTENTIONALLY LEFT BLANK.